

Remarks

Status of application

Claims 1-70 were examined and stand rejected in view of prior art. The claims have been amended to further clarify Applicant's invention. Reexamination and reconsideration are respectfully requested.

The invention

A database system with methodology for parallel schedule generation in a query optimizer is described. In one embodiment, for example, in a database system, a method of the present invention is described for parallel optimization of a query, the method comprises steps of: generating a plurality of parallel plans for obtaining data requested by the query, the parallel plans including parallel operators for executing portions of the query in parallel; adjusting parallel operators of each parallel plan if necessary based on resources available for executing the query; creating a schedule for each parallel plan indicating a sequence for execution of operators of each parallel plan; determining execution cost of each parallel plan based on its schedule; and selecting a particular parallel plan having lowest execution cost for obtaining data requested by the query.

General

A. Sec. 101 rejection

Claims 1-23, 25-68, and 70 stand rejected under 35 U.S.C. 101 on the basis that they are directed to non-statutory subject matter. Here, the Examiner states, for instance, that the claims do not recite a practical application by producing a physical transformation or producing a useful, concrete, and tangible result.

Applicant's independent claim 1 has been amended to set forth the claim limitation of returning a result of a particular parallel plan having lowest execution cost for obtaining data requested by the query. Applicant's other independent claims have been amended in a like manner. It is respectfully submitted that this is a useful, concrete, and tangible result produced by the claim limitations. Accordingly, in view of the amendment, the rejection of Applicant's claims under Section 101 on the basis of non-statutory subject matter is overcome.

B. Rejection under Section 112, second paragraph

Claims 7-10, 13, 34-35, 38, 52, 56, and 59 stand rejected under 35 U.S.C. 112, second paragraph, on the basis of indefiniteness. The claims have been amended to cure any indefiniteness.

Prior art rejections

A. Section 103: Ghosh and Bestgen

Claim 1-70 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Ghosh et al. ("Ghosh", Patent Number 7,051,034) in view of Bestgen et al. ("Bestgen", Patent Number 6,754,652). For the reasons discussed below, Applicant's claimed invention may be distinguished on a variety of grounds.

At the outset, it is important to understand that Applicant does not purport to have invented the broad notion of executing query plans in parallel. A number of different approaches exist for generating plurality of plans, including parallel plans and choosing the best parallel plan.. The notion is therefore well understood in the industry, as well as well documented in the prior art. The focus of Applicant's invention is a specific improvement for optimizing query plans that may execute in parallel, by choosing the best plan. Applicant's way of computing the plurality of different plans and Applicant's way of determining the best parallel plan amongst those plans differ substantially from those prior art approaches. The particular features of this specific improvement are set forth as claim limitations in Applicant's claims and are discussed in further detail below.

The Examiner's argument that Ghosh teaches the art of generating a plurality of plans in a manner similar to Applicant's claimed approach is not supported by Ghosh's own teaching, which limits Ghosh to restartable sub-trees that are applicable only for sub-queries. Here, Ghosh refers to a "re-startable sub-tree," which is a sub-tree created by the query optimizer when a sub-query is found inside of a SQL query. This is a very specialized case (and one which in fact is not even parallelized in Applicant's currently preferred embodiment). The plurality of plans referred to by Ghosh stems from the fact that a sub-tree of this nature could be run in serial or parallel, depending on two aspects of the cost: latency cost and probe cost. Such optimization is deferred until run time to

the query execution engine. Applicant's invention, in contrast, asserts that a plurality of parallel plans is possible and that a schedule may be generated for each of the parallel plans to choose the best possible one based on available resources. This technique is only done at query compilation time. A critical review of Ghosh reveals that the reference does not describe a way to generate multiple parallel plans and reveals that the reference discloses nothing about how to choose the most optimal plan. Instead, Ghosh's teaching simply focuses on a new way for processing sub-queries.

Bestgen, for its part, does provide a way to generate plurality of plans, but the invention is limited to choosing the best plan in an extensible optimizer framework. Bestgen shows that a plurality of plans is a necessary step in the query optimization process and provides a flexible framework to incorporate newly discovered knowledge into this framework using an objected oriented rule driven framework. However, Bestgen reveals nothing about (i) the methodology to generate a plurality of parallel plans, (ii) adjusting parallel plans based on available resources, and (iii) the way to choose the best parallel plan. Similarly, Bestgen use of cost metrics differs substantially from Applicant's use. Applicant's invention, in contrast, generates a schedule for each potential parallel plan, with the best plan being chosen based on the execution cost of each schedule.

In Applicant's claimed invention, parallel operators of each parallel plan are adjusted if necessary based on resources available for executing the query (see e.g., amended claim 1). Parallel operators refer to constructs that reflect SQL operations, such as join operations, predicate evaluation mechanisms, scans (e.g., table or index), grouping operations, union operations, distinct operations -- that is, operations defined by SQL. In Applicant's system, these operations are examined to find a way to parallelize those operations. However, the particular focus of Applicant's invention is finding out the best possible parallel plan. This is done by looking at available resources, costing those operators against the resources actually available, creating possible plans (i.e., each plan associated with a schedule indicating how the operators are going to be activated, and how cloned operators are going to be activated for parallel operation), determining the response time for the possible schedules, and finally selecting the one plan that is the cheapest (i.e., optimal). The creation of a schedule for each parallel plan (i.e., indicating

a sequence for execution of operators of each parallel plan and associated costs therein) is absolutely novel, and is not taught or suggested by the cited art.

In order to appreciate these features of Applicant's invention, it is perhaps instructive to examine how Applicant's system processes queries. Consider the following sample query:

```
select distinct CUST_ID
from CUSTOMER
group by STATE
```

From the SQL perspective for this query, the database system will scan the CUSTOMER table, perform some sort of distinct analysis on CUST_ID (which itself may already be unique) followed by grouping/aggregation, and thereafter provide the results. From the perspective of Applicant's invention, the underlying operations may be performed in a multitude of different ways. For example, the system could scan the whole table serially or in parallel. When performing the distinct analysis, the system could do that in parallel if the scan is performed in parallel. In that case, as the scan is being performed in parallel, the result set from the scan could be merged into one stream that is funneled into the distinct operation or re-split the data streams into a larger number of sub-streams so that the distinct operation can be done with a greater degree of parallelism. Similarly, for the grouping operation, the system could do that serially or in parallel, and so forth and so on.

As can be seen, for each operation Applicant's system may decide to do it serially or in parallel. Given that, there is a wide range of possibilities available in how serial and parallel operations (including different levels of parallelism within a given operation) are combined. There is another dimension of choices is available, which is a scan operation may be performed using a table scan or an index scan (as available). For a distinct operation, the system may perform a hash-based distinct operation or sort-based distinct operation. Similarly, for grouping, the system could do a hash-based grouping or sort-based grouping, or if the data is already sorted within the system, it could perform the grouping on the already-sorted data. As should be readily apparent, combining these two dimensions of possibilities in all possible ways yields n different plans, where n may be a very large number.

For all of those possible plans, the system may compute a cost to determine which one is the winner (i.e., best plan). To figure out the cost (elapsed time) for each possible plan, Applicant's system figures out how all the various operators are to be scheduled. In other words, when exactly would an operator start and when would it stop, based on available resources. That information comprises the activation schedule for a given operation. Based on that information, that is the activation schedule of various operators, Applicant's system can determine the elapsed time for each particular plan, and thus can determine which plan is in fact the best one.

All told, none of the prior art cited by the Examiner shows anything remotely close to Applicant's schedule-based approach. Ghosh, for example, lacks any notion of a schedule. Even his costing fails to take into account available resources for operations performed in parallel. Instead, he focuses on the cost of doing a plan in serial (e.g., the I/O time for doing a plan serially) and determining what is the threshold for performing the plan in parallel (i.e., makes parallel attractive). Bestgen, which itself fails to cure these deficiencies of Ghosh, provides nothing that teaches or suggests Applicant's schedule-based approach. Accordingly, in view of the above remarks (as well as clarifying amendments made to Applicant's independent claims), it is respectfully submitted that the claims set forth a patentable advance over the art, and that any rejection under Section 103 is overcome.

Any dependent claims not explicitly discussed are believed to be allowable by virtue of dependency from Applicant's independent claims, as discussed in detail above.


Conclusion

In view of the foregoing remarks and the amendment to the claims, it is believed that all claims are now in condition for allowance. Hence, it is respectfully requested that the application be passed to issue at an early date.

If for any reason the Examiner feels that a telephone conference would in any way expedite prosecution of the subject application, the Examiner is invited to telephone the undersigned at 408 884 1507.

Respectfully submitted,

Date: January 3, 2007



Digitally signed by John A. Smart
DN: cn=John A. Smart, o, ou,
email=JohnSmart@Smart-IPLaw.
com, c=US
Date: 2007.01.03 18:53:14 -08'00'

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX